

WHAT IS CLAIMED IS:

1        ✓ 1. A method for local area network (LAN) emulation over an Infiniband (IB) fabric  
2 comprising:

3            receiving, at an IB LAN driver at a first node on an IB fabric, at least one port and  
4 associated local identifier (LID) of at least one remote peer node on the IB fabric;

5            generating an IEEE 802.3 Ethernet Media Access Control (MAC) address with  
6 one LID imbedded, the imbedded LID being for at least one remote peer node, the IB  
7 LAN driver sending the Ethernet MAC address to an Address Resolution Protocol  
8 (ARP);

9            generating a logical address of the at least one remote peer node by a network  
10 protocol;

11           mapping the logical address to the Ethernet MAC address;

12           sending, by the IB LAN driver, the Ethernet MAC address onto the IB fabric to  
13 the at least one remote peer node, the at least one remote peer node appearing to  
14 reside on an Ethernet network to the network protocol.

1           2. The method according to claim 1, wherein the port and LID of the at least one  
2 remote peer node is received from a name service.

1           3. The method according to claim 1, wherein the name service receives the port  
2 and LID of the at least one remote peer node from a subnet manager in the IB fabric.

1           4. The method according to claim 1, further comprising performing the mapping  
2 of the logical address to the physical address by an Address Resolution Protocol (ARP).

1           5. The method according to claim 1, wherein the network protocol comprises  
2 one of NetWare, Open Systems Interconnection (OSI), Transmission Control  
3 Protocol/Internet Protocol (TCP/IP), DECnet, and AppleTalk.

1           6. The method according to claim 1, further comprising mapping the LID into the  
2 least significant sixteen bits of the Ethernet MAC address.

1           7. The method according to claim 1, wherein the Ethernet MAC address  
2 comprises a broadcast address to all at least one remote peer nodes.

1           8. The method according to claim 1, wherein the Ethernet MAC address  
2 comprises a multicast address to some of the at least one remote peer nodes.

1           9. The method according to claim 1, wherein the Ethernet MAC address  
2 comprises a unicast address to one of the at least one remote peer nodes.

1           / 10. A node on an Infiniband (IB) fabric comprising:  
2 a channel adapter containing at least one port providing access to the IB fabric,  
3 each port having a local identifier (LID);

4 a name service, the name service obtaining at least one port and at least one  
5 LID for at least one remote peer node on the IB fabric;

6 at least one network protocol, the at least one network protocol generating a  
7 logical address of the at least one remote peer node to send data;

8 an Address Resolution Protocol (ARP), the ARP mapping the logical address to  
9 a physical address, the physical address being an IEEE 802.3 Ethernet Media Access  
10 Control (MAC) address imbedded with the LID of the at least one remote peer node;  
11 and

12 an IB local area network (LAN) driver, the IB LAN driver providing unicast,  
13 multicast, and broadcast capability for transfers across the IB fabric to the at least one  
14 remote peer node, the IB LAN driver sending the Ethernet MAC address and the data  
15 to the at least one remote peer node through at least one port, the at least one remote  
16 peer node appearing to reside on an Ethernet network to the network protocol.

17 11. The node according to claim 10, further comprising a transport services  
18 library (TSL), the TSL providing connection management, work queue management,  
19 memory management, and message pool management, the IB LAN driver using the  
20 TSL to establish a connection with and perform transfers to the at least one remote  
21 peer node.

22 12. The node according to claim 10, further comprising an IB bus driver, the IB  
23 bus driver loading the IB LAN driver at the node when the at least one port of the  
24 channel adapter is initialized and set active, the IB bus driver receiving each LID and

4 a LID mask for each LID from the IB LAN driver once the port is activated and assigning  
5 one LID to each at least one port.

1 13. The node according to claim 12, the IB bus driver using a vendorID and a  
2 deviceID to locate and load the appropriate IB LAN driver on the node.

1 14. The node according to claim 12, the at least one port of the channel adapter  
2 being initialized and set active by a subnet manager on the IB fabric.

3 15. The node according to claim 12, wherein the name service obtains the at  
4 least one port and the at least one LID for the at least one remote peer node on the IB  
5 fabric from a Subnet Management Database (SMDB), the SMDB residing on the IB  
6 fabric and providing persistent storage of subnet topology, subnet events, and subnet  
7 configuration information.

1 16. The node according to claim 11, wherein the maximum transmission unit  
2 (MTU) of the IB LAN driver is configurable and is set larger than the maximum packet  
3 size allowed on the IB fabric.

1 17. The node according to claim 16, wherein the TSL receives the data and  
2 segments the data into a packet size compatible with the IB fabric.

1 18. The node according to claim 11, the TSL further comprising a queue pair for  
2 each connection between the node and one at least one remote peer node, only one  
3 queue pair being used for broadcast transfers to all at least one remote peer node.

1 / 19. An article comprising a storage medium with instructions stored therein, the  
2 instructions when executed causing a processing device to perform:

3 receiving a port and a local identifier (LID) of a local node on an Infiniband (IB)  
4 fabric;

5 generating an IEEE 802.3 Ethernet Media Access Control (MAC) address with  
6 the LID imbedded and sending the Ethernet MAC address to an Address Resolution  
7 Protocol (ARP);

8 receiving at least one port and associated local identifier (LID) for at least one  
9 remote peer node on the Infiniband (IB) fabric;

10 generating at least one second Ethernet MAC address with the LID of the at least  
11 one remote peer node imbedded and sending the at least one second Ethernet MAC  
12 address to the Address Resolution Protocol (ARP);

13 sending at least one second Ethernet MAC address onto the IB fabric to at least  
14 one remote peer node in response to a network protocol request, the at least one  
15 remote peer node appearing to reside on an Ethernet network to the network protocol.

1 20. The article according to claim 19, wherein the Ethernet MAC address  
2 comprises 48 bits.

1

1